Recovering Metabolic Networks using A Novel Hyperlink Prediction Method

Muhan Zhang[1], Zhicheng Cui[1], Tolutola Oyetunde[2], Yinjie Tang[2], and **Yixin Chen[1]**
[1]Department of Computer Science and Engineering, Washington University in St. Louis
[2]Department of Energy, Environmental & Chemical Engineering, Washington University in St. Louis
{muhan, z.cui, toyetunde, yinjie.tang}@wustl.edu, chen@cse.wustl.edu

arXiv:1610.06941v2 [cs.SI] 21 Nov 2016

Abstract

Studying metabolic networks is vital for many areas such as novel drugs and bio-fuels. For biologists, a key challenge is that many reactions are impractical or expensive to be found through experiments. Our task is to recover the missing reactions. By exploiting the problem structure, we model reaction recovery as a hyperlink prediction problem, where each reaction is regarded as a hyperlink connecting its participating vertices (metabolites). Different from the traditional link prediction problem where two nodes form a link, a hyperlink can involve an arbitrary number of nodes. Since the cardinality of a hyperlink is variable, existing classifiers based on a fixed number of input features become infeasible. Traditional methods, such as common neighbors and Katz index, are not applicable either, since they are restricted to pairwise similarities. In this paper, we propose a novel hyperlink prediction algorithm, called Matrix Boosting (MATBoost). MATBoost conducts inference jointly in the incidence space and adjacency space by performing an iterative completion-matching optimization. We carry out extensive experiments to show that MATBoost achieves state-of-the-art performance. For a metabolic network with 1805 metabolites and 2583 reactions, our algorithm can successfully recover nearly 200 reactions out of 400 missing reactions.

Introduction

Reconstructed metabolic networks are important tools for understanding the metabolic basis of human diseases, increasing the yield of biologically engineered systems, and discovering novel drug targets (Bordbar et al. 2014). Semi-automated procedures have been recently developed to reconstruct metabolic networks from annotated genome sequences (Thiele and Palsson 2010). However, these networks are often incomplete – some reactions can be missing from them, which can severely impair their utility (Kumar, Dasika, and Maranas 2007). Thus, it is critical to develop accurate computational methods for completing metabolic networks. A method that can automatically recover missing reactions in metabolic networks is much desired.

Link prediction (Liben-Nowell and Kleinberg 2007; Lü and Zhou 2011) has been studied broadly in recent years (Singh and Gordon 2008; Miller, Jordan, and Griffiths 2009; Rendle et al. 2009; Menon and Elkan 2011; Cao, Kong, and Yu 2014; Chen et al. 2015; Ermiş, Acar, and Cemgil 2015; Song, Meyer, and Tao 2015; Wu et al. 2016). Existing methods can be grouped into three types: topological feature-based approaches, latent feature-based approaches, and supervised learning approaches. Popular criterions and methods include common neighbors, Jaccard coefficients (Liben-Nowell and Kleinberg 2007), Katz index (Katz 1953), and latent feature models (Miller, Jordan, and Griffiths 2009; Rendle et al. 2009; Menon and Elkan 2011) et al. However, these approaches are restricted to predicting pairwise relations. None of them are directly applicable to predicting hyperlinks. A hyperlink relaxes the assumption that only two nodes can participate in a link. Instead, an arbitrary number of nodes are allowed to jointly form a hyperlink. A network made up of hyperlinks is called a hypernetwork or hypergraph.

Many complex networks are actually hypernetworks. Examples include biological networks, social networks with group relations, and citation networks. In metabolic networks, every reaction can be regarded as a hyperlink among its component metabolites. A hypernetwork can be naturally represented as an *incidence matrix* S, where each column of S is a hyperlink and each row represents a vertex. If vertex i participates in hyperlink j, then $S_{ij} = 1$, otherwise $S_{ij} = 0$. In this paper, we model metabolic reaction recovery as a hyperlink prediction problem in hypernetworks. For metabolic networks, the S matrix is also known as the *Binary Stoichiometric Matrix*. Figure 1 shows an example.

Although hypernetworks are common in the real world, there are still limited research on hyperlink prediction. One great challenge lies in the variable cardinality of hyperlinks. Existing classifiers for link prediction are based on a fixed number of input features (e.g., concatenating features of vertex a and features of vertex b). However, in hyperlink prediction, the number of vertices in a hyperlink is variable, which makes existing methods inapplicable. On the other hand, link prediction methods based on topological features, such as common neighbors, cannot be applied to hyperlink prediction either. The reason is that these measures are defined for pairs of nodes (links) instead of hyperlinks. A few basic methods for hyperlink prediction have poor performance on recovering metabolic networks. In summary, hyperlink prediction is an *important, interesting, yet rarely touched challenging problem*.

In this paper, we introduce a novel hyperlink prediction algorithm for recovering metabolic networks. Our key ob-

servation is that a hyperlink $\mathbf{s}$ (a column vector in an incidence matrix) can be transformed into its equivalent matrix expression in the vertex adjacency space by $\mathbf{s}\mathbf{s}^{\mathrm{T}}$. In this way, we are able to first infer the pairwise relationships between vertices leveraging existing link prediction methods, and then recover the missing hyperlinks through constrained optimization. Based upon this, we propose a two-step EM-like optimization method, which alternately performs a C step (completion) and an M step (matching) to find the best hyperlinks. We call the resulting algorithm Matrix Boosting (MATBoost). We compare our algorithm with six baseline methods on eight metabolic networks from five species. Experimental results demonstrate that our algorithm outperforms the state-of-the-art methods by a large margin on recovering missing metabolic reactions.

Reaction Recovery as Hyperlink Prediction

One significant difference between link prediction and hyperlink prediction is the total number of possible links / hyperlinks. For a network with M vertices, the total number of links is $O(M^2)$, while the total number of hyperlinks will be $O(2^M)$. This prevents us from giving a score to every possible hyperlink like what we do in link prediction. Therefore, we consider a *transductive inference* setting, where a set of test hyperlinks are given as candidates, and we only classify the given test examples instead of classifying all. This setting is natural in practical applications. For instance, in the reaction prediction task, we do not need to consider all 2^M possible reactions since most of them do not contain biological meanings. Instead, we restrict the candidate hyperlinks to be within a universal reaction pool. This pool is the set of all known metabolic reactions. Consequently, our task is to classify which reactions in the pool are the missing reactions for a specific metabolic network.

Problem definition and notations

We formally define the missing reaction recovery problem as follows: Given an $M \times N^S$ binary stoichiometric matrix S (M metabolites and N^S reactions) of a metabolic network, we assume that several reactions are missing from it. We denote these missing columns by ΔS. We also have a universal reaction pool, whose binary stoichiometric matrix is denoted by U. U contains all known metabolic reactions, including both positive reactions (ΔS) as well as negative reactions which do not belong to this metabolic network. Our task is to recover/predict as many missing reactions from U as possible. The corresponding hyperlink prediction problem is parallel. Given a training incidence matrix S (where some column hyperlinks ΔS are missing) and a testing incidence matrix U (the candidate hyperlinks to select from), we aim to predict which columns in U belong to ΔS.

For an incidence matrix S, we can project it into the vertex adjacency space by $A = SS^{\mathrm{T}}$. A is the weighted adjacency matrix where A_{ij} is the number of hyperlinks that vertices i and j both participate in. In systems biology, A is called the *Compound Adjacency Matrix*.

Matrix Boosting Algorithm

We illustrate the proposed Matrix Boosting (MATBoost) algorithm in Figure 1. The overall framework is as follows.

Input: Initial training adjacency matrix $A = SS^{\mathrm{T}}$, candidate hyperlinks U, initial zero prediction scores Λ_0, maximum iteration number $I_{\max}$.

Repeat from $k = 1$:

a) Update A, set:

$$A_k = A + [U\Lambda_{k-1}U^{\mathrm{T}}]_A. \tag{1}$$

b) The C (completion) step: Input A_k, output the matrix completion results $\Delta\hat{A}_k$.

c) The M (matching) step: Input $\Delta\hat{A}_k$, output the prediction scores Λ_k.

d) Stopping criterion: If $||\Lambda_k - \Lambda_{k-1}||_F \geq ||\Lambda_{k-1} - \Lambda_{k-2}||_F$ or $k = I_{\max}$, stop the algorithm; otherwise $k \leftarrow k + 1$ and repeat from step a).

Output: The final prediction scores:

$$\bar{\Lambda} = \frac{1}{k-2}\sum_{i=1}^{k-2}\Lambda_i. \tag{2}$$

We now explain each step in the following subsections.

Building block 1: decomposition of the adjacency matrix

Since direct inference in the incidence space is difficult (due to the variable cardinality of hyperlinks), we first project hyperlinks into their vertex adjacency space and infer their vertex interactions. Given the training hyperlinks S, we can calculate the training adjacency matrix by $A = SS^{\mathrm{T}}$. However, since S is not complete (some columns ΔS are missing), the resulting A is also incomplete.

Because ΔS is missing, the complete incidence matrix should be $[\,S\ \Delta S\,]$. We can calculate its adjacency matrix as follows:

$$\begin{aligned}[\,S\ \Delta S\,]\,[\,S\ \Delta S\,]^{\mathrm{T}} &= SS^{\mathrm{T}} + \Delta S\Delta S^{\mathrm{T}} \\ &= A + \Delta A,\end{aligned} \tag{3}$$

where we define $\Delta A = \Delta S\Delta S^{\mathrm{T}}$. We notice that since hyperlinks ΔS are missing, the adjacency matrix A also incurs a loss ΔA. If we can accurately recover ΔA, we will be able to infer ΔS consequently.

A straightforward idea is to perform link prediction on ΔA using A. However, this will cause problems because ΔA may overlap with A at some entries, while existing link prediction methods require that training links A and testing links ΔA are fully separated.

To fix this issue, we first decompose $A + \Delta A$ as follows:

$$A + \Delta A = \underbrace{A + [\Delta A]_A}_{A^+} + \underbrace{[\Delta A]_{\bar{A}}}_{\Delta A^-}, \tag{4}$$

where we have used $[X]_A$ to denote the operation that only keeps the entries of X at A's nonempty entries and mask all else, which is defined as:

$$[X_{ij}]_A = \begin{cases} X_{ij}, & \text{if } A_{ij} \neq 0 \\ 0, & \text{if } A_{ij} = 0. \end{cases} \tag{5}$$

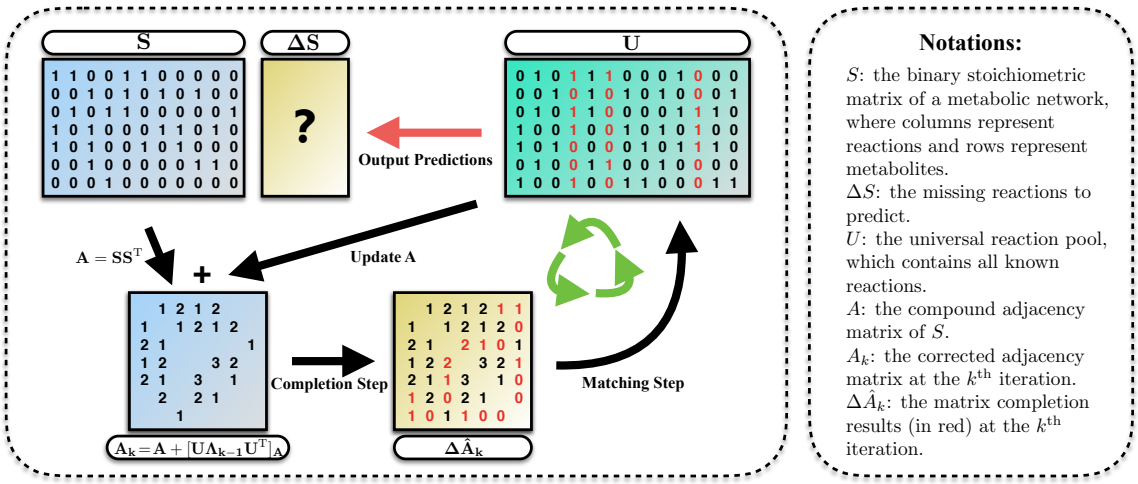


Figure 1: An illustration of MATBoost. The training incidence matrix S is first transformed to its adjacency matrix A. A has many missing entries some of which are from ΔS. The C (completion) step uses matrix factorization to complete the missing entries. The M (matching) step will search for the best matches from the candidate matrix U by solving a constrained optimization. The matched hyperlinks will then be added to A in order to correct A's bias. This procedure is iterated several times to boost the performance.

$[X]_{\bar{A}}$ is conversely defined as keeping X only at A's empty entries.

Equation (4) decomposes the complete adjacency matrix into A^+ and ΔA^-. A^+ is the training adjacency matrix A plus the entries of ΔA that overlap with A. ΔA^- is ΔA discarding its entries that overlap with A. Such a decomposition ensures that A^+ and ΔA^- are isolated from each other. Our goal is thus to leverage A^+ to predict ΔA^-, and subsequently, to recover ΔS.

Building block 2: the completion step

We first assume that A^+ is known (which is not true because ΔA is unknown beforehand, we will show how to iteratively approximate it later). Then predicting ΔA^- can be regarded as a matrix completion (Candès and Recht 2009) or weighted link prediction problem (Lü and Zhou 2010). Concretely, we aim to complete the empty entries of A^+, which we denote by $\Delta\hat{A}$. Ideally, the predicted $\Delta\hat{A}$ will resemble ΔA^- by having large values at ΔA^-'s nonempty entries and near-zero values at other places.

In this paper, we adopt matrix factorization (Koren, Bell, and Volinsky 2009) to infer ΔA^- because it is fast and accurate. The optimization problem is as follows:

$$\begin{aligned} &\underset{\Theta}{\text{minimize}} \sum_{i<j, A^+_{ij}>0} ||A^+_{ij} - y_{ij}||_2^2 + \gamma\mathcal{R}(\Theta), \\ &\text{where } y_{ij} = w_0 + w_i + w_j + \sum_{f=1}^{k} v_{if}v_{jf}, \\ &\Theta \text{ is the set of all parameters } w_0, w_i, w_j, v_{if}, v_{jf}. \end{aligned} \tag{6}$$

The optimization only uses half of the training entries since the adjacency matrix A^+ is symmetric. After training, we can complete the empty entries by:

$$\Delta\hat{A}_{ij} = \begin{cases} w_0 + w_i + w_j + \mathbf{v}_i^{\mathrm{T}}\mathbf{v}_j, & \text{if } A^+_{ij} = 0, \\ 0, & \text{if } A^+_{ij} \neq 0. \end{cases} \tag{7}$$

Building block 3: the matching step

In the matching step, we aim to select a a batch of hyperlinks from U that best match the predicted $\Delta\hat{A}$. Firstly, we define a diagonal matrix Λ containing the selection indicators λ_i for the i^{th} hyperlink in U. λ_i is 1 if hyperlink i is selected and 0 otherwise.

Now consider the optimization problem as follows:

$$\begin{aligned} &\underset{\Lambda}{\text{minimize}} \;\; ||[U\Lambda U^{\mathrm{T}}]_{\bar{A}} - \Delta A^-||_F^2, \\ &\text{subject to } \;\lambda_i \in \{0,1\},\; i = 1 \ldots N^U, \\ &\text{where } \;\Lambda = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ldots & \\ & & & \lambda_{N^U} \end{bmatrix}, \\ &N^U \text{ is the number of columns in } U, \\ &||\cdot||_F \text{ is the Frobenius norm.} \end{aligned} \tag{8}$$

Theorem 1. *The optimization problem (8) is an integer least square problem with the ground-truth indicators Λ^* being a global minimum of it. (proof in the supplemental material)*

In practice, we will substitute $\Delta\hat{A}$ for ΔA^- in (8) to solve Λ. However, integer least square (ILSQ) optimization is NP hard. When N^U is large (usually thousands in metabolic networks), it is generally intractable. Therefore, we propose a constrained lasso regression instead. The optimiza-

tion problem is as follows:

$$\begin{aligned} &\underset{\Lambda}{\text{minimize}} \;\; ||[U\Lambda U^{\mathrm{T}}]_{\bar{A}} - \Delta\hat{A}||_F^2 + \alpha \sum_{i=1}^{N^U} |\lambda_i|, \\ &\text{subject to} \;\; 0 \le \lambda_i \le 1, \; i = 1 \; ... \; N^U, \\ &\text{where} \;\; \Lambda = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & ... & \\ & & & \lambda_{N^U} \end{bmatrix}, \end{aligned} \tag{9}$$

The above optimization relaxes the integer λ_i to be continuous within $[0, 1]$. Because of the limited number of hyperlinks in U that contribute to ΔA^-, we add an lasso (L_1) penalty to preserve the sparsity of solutions. We solve the above optimization using subgradient methods. After getting the prediction scores λ_i, we can rank the hyperlinks in U and select the top ones as our predictions of the missing hyperlinks.

Combine all: boosting via iterative optimization

In the completion step, we have assumed that $A^+ = A + [\Delta A]_A$ is known in order to separate ΔA^- from A^+ and infer ΔA^-. Notice that A^+ is essentially the training adjacency matrix A plus the overlapping entries of the missing hyperlinks $\Delta S \Delta S^{\mathrm{T}}$. Thus, to approximate the unknown A^+, we propose to iteratively add the predicted missing hyperlinks $U\Lambda U^{\mathrm{T}}$ into A (Λ is real matrix now) to remedy its bias from the real A^+.

Given the last iteration's prediction scores Λ_{k-1}, we first update A_k by adding to A the predicted hyperlinks $[U\Lambda_{k-1}U^{\mathrm{T}}]_A$ from the last iteration. After getting A_k, which is an approximation of A^+, the algorithm performs the C step and M step subsequently. MATBoost performs this procedure in an iterative fashion and outputs the average prediction scores in order to boost the performance. It is further equipped with an early stopping criterion which we will explain in the next section.

Algorithm Analysis

We provide two different views of the MATBoost algorithm in this section. First we define some notations:

Definition 1. *We define the C (completion) step (6) and (7) as a map $T_c : \mathcal{A} \to \mathcal{A}$, where $\mathcal{A}$ is the space of adjacency matrices A. T_c takes an incomplete adjacency matrix A_k as input and outputs the predictions $\Delta\hat{A}_k$.*

Definition 2. *We define the M (matching) step (9) as a map $T_m : \mathcal{A} \to \mathcal{L}$, where $\mathcal{L}$ is the space of diagonal prediction score matrices Λ. T_m takes the predictions $\Delta\hat{A}_k$ as input and outputs the prediction scores Λ_k.*

A coordinate descent view

MATBoost may be viewed as an example of coordinate descent algorithms. In the C step, we fix Λ and update $\Delta\hat{A}$. In the M step, we fix $\Delta\hat{A}$ and update Λ. Concretely, MATBoost can be viewed as an alteration between the two steps:

C (completion) step:

Fix Λ_{k-1}, update $\Delta\hat{A}$ by $\Delta\hat{A}_k = T_c(A + [U\Lambda_{k-1}U^{\mathrm{T}}]_A)$.

M (matching) step:

Fix $\Delta\hat{A}_k$, update Λ by $\Lambda_k = T_m(\Delta\hat{A}_k)$.

A fixed-point view

Here we discuss the convergence properties and the stopping criterion of MATBoost under a fixed-point view.

Definition 3. *We define one iteration of both C and M steps as a map $T_{cm} : \mathcal{L} \to \mathcal{L}$. T_{cm} takes the Λ_{k-1} from the last iteration as input and outputs the new prediction scores Λ_k. Equivalently, $T_{cm}(\Lambda_{k-1}) = T_m(T_c(A + [U\Lambda_{k-1}U^{\mathrm{T}}]_A))$.*

Having defined T_{cm}, we find that the MATBoost algorithm can be expressed as a fixed-point iteration algorithm (Khamsi and Kirk 2011), where the algorithm iteratively takes the intermediate result $\Lambda_k = T_{cm}(\Lambda_{k-1})$ as input to get $\Lambda_{k+1}, \Lambda_{k+2}$... until the stopping criterion is satisfied.

Now we introduce the Banach fixed-point theorem.

Theorem 2. *(Banach fixed-point theorem (Palais 2007)) If T_{cm} is a contraction mapping, then MATBoost will iteratively converge to a fixed-point Λ^*, which satisfies $\Lambda^* = T_{cm}(\Lambda^*)$.*

Theorem 3. *If T_{cm} is a contraction mapping defined on the metric space $(\mathcal{L}, d)$ and $d(\Lambda_x, \Lambda_y) = ||\Lambda_x - \Lambda_y||_F$, then the intermediate results $\Lambda_1, ..., \Lambda_{k-2}$ of MATBoost satisfy the following inequality for any $k > 2$:*

$$||\Lambda_{k-1} - \Lambda_{k-2}||_F < ||\Lambda_{l-1} - \Lambda_{l-2}||_F, \;\; \forall\, 2 \le l < k. \tag{10}$$

Besides, Λ_{k-2} is the closest point to Λ^* among $\Lambda_0, ..., \Lambda_{k-2}$. (proof in the supplemental material)

From Theorem 3 we see that $||\Lambda_k - \Lambda_{k-1}||_F < ||\Lambda_{k-1} - \Lambda_{k-2}||_F$ for all k is a necessary condition for T_{cm} being a contraction mapping on $(\mathcal{L}, d)$. Recall that the stopping criterion in MATBoost is $||\Lambda_k - \Lambda_{k-1}||_F \ge ||\Lambda_{k-1} - \Lambda_{k-2}||_F$. We can see that it is essentially tracking this necessary condition. Once it is broken, the fixed-point iteration will stop converging. And the algorithm will stop the iteration earlier by outputting the average scores before iteration $k - 1$.

Since $||T_{cm}(\Lambda^*) - \Lambda^*||_F = 0$, we may also regard $||\Lambda_k - \Lambda_{k-1}||_F = ||T_{cm}(\Lambda_{k-1}) - \Lambda_{k-1}||_F$ as a measure of how close the iteration $k - 1$ is to convergence.

In practice, we find that $||\Lambda_k - \Lambda_{k-1}||_F$ typically drops sharply in the first iteration, then keeps decreasing for a few rounds, and then starts to fluctuate around a small value. Instead of using the last Λ_{k-2}, we find that using the mean scores of Λ_1 to Λ_{k-2} usually achieves higher performance. We have Theorem 4 to explain it.

Theorem 4. *The MATBoost algorithm is an ensemble of $k-2$ weak learners of T_{cm} with different input matrices $A + [U\Lambda_i U^{\mathrm{T}}]_A$, i from 0 to $k - 3$.*

Proof. It is easy to prove by noticing that Λ_1 to Λ_{k-2} are just the outputs of $T_{cm}(\Lambda_0)$ to $T_{cm}(\Lambda_{k-3})$. Thus the final prediction scores can be written as:

$$\bar{\Lambda} = \frac{1}{k-2}\sum_{i=1}^{k-2}\Lambda_i = \frac{1}{k-2}\sum_{i=0}^{k-3}T_{cm}(\Lambda_i). \tag{11}$$

The theorem follows. □

Remark: We may regard $A+[U\Lambda_i U^{\mathrm{T}}]_A$ as noisy samplings of A^+. Therefore, MATBoost can benefit from ensemble learning which is expected to improve performance.

Related Work

Although hyperlinks are common in real world and can be used to model multi-way relationships, currently there are still limited research on hyperlink prediction. Xu et al. (Xu, Rockmore, and Kleinbaum 2013) proposed a supervised HPLSF framework to predict hyperlinks in social networks. To deal with the variable number of features, HPLSF uses their entropy score as a fixed-length feature for training a classification model. To our best knowledge, this is the only algorithm that is specifically designed for hyperlink prediction in arbitrary-cardinality hypernetworks.

Nevertheless, learning with hypergraphs as a special data structure is popular in the machine learning community, e.g., semi-supervised learning with hypergraph regularization (Yu, Tao, and Wang 2012; Tian, Hwang, and Kuang 2009), modeling label correlations via hypernetworks in multi-label learning (Sun, Ji, and Ye 2008), and modeling communities to improve recommender systems (Bu et al. 2010). Zhou et al. (Zhou, Huang, and Schölkopf 2006) studied spectral clustering in hypergraphs. They generalized the normalized cut (Shi and Malik 2000) algorithm to hypergraph clustering and proposed a hypergraph Laplacian. They also adapted their previous semi-supervised learning algorithm using graph regularization (Zhou, Huang, and Schölkopf 2005) to hypergraph vertex classification. These research mainly aim to improve the learning performance on nodes by leveraging their hyperlink relations. However, none of them focuses on predicting the hyperlink relations.

Experimental Results

To demonstrate the effectiveness of the proposed Matrix Boosting (MATBoost) algorithm in reaction recovery tasks, we compare it with six baseline methods on eight metabolic networks. The datasets and source code in this paper are available online: https://github.com/ufo0010/HyperLinkPrediction.

Baselline algorithms

We compare the proposed Matrix Boosting algorithm with the following six baseline algorithms:

SHC (Spectral Hypergraph Clustering) (Zhou, Huang, and Schölkopf 2006) is a state-of-the-art hypergraph learning algorithm. SHC outputs classification scores by $f = (I - \xi\Theta)^{-1}y$. The hyperparameter ξ is determined by searching over the grid $\{0.01,0.1,0.5,0.99,1\}$ using five-fold cross validation. SHC is originally designed to classify hypergraph vertices leveraging their hyperlink relations. Here we transpose the incidence matrices to change each vertex into a hyperlink and each hyperlink into a vertex. Thus SHC can be used for hyperlink prediction.

HPLSF (Xu, Rockmore, and Kleinbaum 2013) is a hyperlink prediction method using supervised learning. It calculates an entropy score among vertices along each latent feature dimension in order to get a fixed-length feature input. We train a logistic regression model on these entropy features in order to output prediction scores.

FM (Factorization Machine) (Rendle 2010; Rendle 2012) is a flexible factorization model. We use the classification function of FM, where columns of the training incidence matrix are directly used as input features to the model.

HKatz is a generalization of Katz index (Katz 1953) for hyperlinks. Concretely, a hyperlink containing m vertices will have $m(m-1)/2$ pairwise Katz indices. We calculate their average as the HKatz (Hypernetwork Katz) index of the hyperlink. The parameter β is determined by searching over $\{0.001,0.005,0.01,0.1,0.5\}$ using five-fold cross validation.

HCN is a generalization of common neighbors (Liben-Nowell and Kleinberg 2007) for hyperlinks. For a hyperlink containing m vertices, the average of the $m(m-1)/2$ pairwise common neighbors is calculated as the HCN (Hypernetwork Common Neighbors) score for this hyperlink.

Random: a theoretical baseline for comparing algorithms' performance against random. It is equal to assigning random scores between [0,1] to all testing hyperlinks.

In MATBoost and FM, libFM (Rendle 2012) is adopted as the default matrix factorization software. The optimization scheme is MCMC, which has advantages of high accuracy and automatic regularization. The latent factor number k is set to 8 (default) in both algorithms. The lasso hyperparameter α and the maximum iteration number of MATBoost is set to 0.1 and 10 respectively.

Evaluation and implementation

We evaluate the reaction recovery performance based on their prediction scores. We adopt AUC (area under the ROC curve) as one measure. We also evaluate performance based on how many missing reactions can be successfully recovered if we select top n reactions in U as predictions (n is the number of missing reactions). We call the second measure "Recovered number". Compared to AUC, this measure only focuses on top predictions and can more directly reflect how well hyperlink prediction algorithms perform in practice.

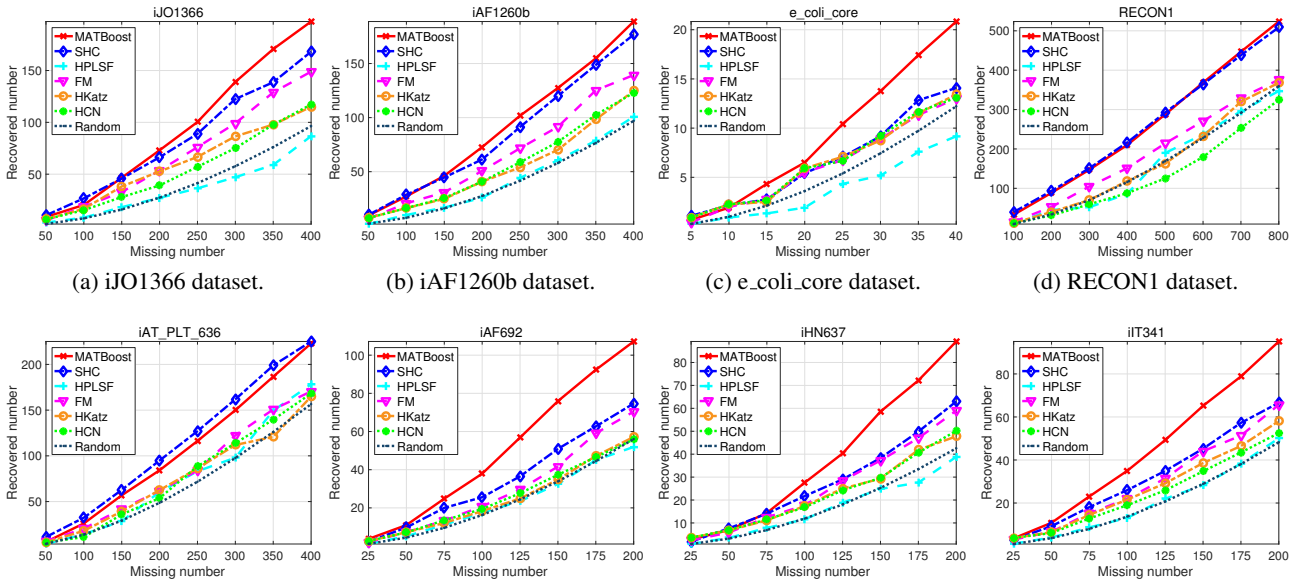
We implement our algorithms mainly in MATLAB. All experiments are done on a twelve-core Intel Xeon Linux server. All experiments are repeated twelve times and the average results are presented.

Datasets and results

We download all 11893 reactions from BIGG[1] to build a universal reaction pool. These reactions are collected from 79 metabolic networks of various organisms. We conduct experiments on eight metabolic networks from five species: E. coli, H. sapiens, M. barkeri, Cl. ljungdahlii, and H. pylori. For each metabolic network, we filter out the universal reactions which contain exotic metabolites from the pool. We also filter out the network's existing reactions from the pool. Table 1 shows the statistics of the eight datasets.

For each dataset, we randomly delete [5:5:40], [25:25:200], [50:50:400] or [100:100:800] (according to network size) reactions as missing hyperlinks, and keep

[1] http://bigg.ucsd.edu



(a) iJO1366 dataset. (b) iAF1260b dataset. (c) e_coli_core dataset. (d) RECON1 dataset.

(e) iAT_PLT_636 dataset. (f) iAF692 dataset. (g) iHN637 dataset. (h) iIT341 dataset.

Figure 2: Number of recovered reactions for different numbers of missing reactions.

Dataset	Species	Vertices	Hyperlinks	Neg. Hyperlinks
(a) iJO1366	E. coli	1805	2583	1253
(b) iAF1260b	E. coli	1668	2388	1242
(c) e_coli_core	E. coli	72	95	91
(d) RECON1	H. sapiens	2766	3742	973
(e) iAT_PLT_636	H. sapiens	738	1008	621
(f) iAF692	M. barkeri	628	690	513
(g) iHN637	Cl. Ljungdahlii	698	785	740
(h) iIT341	H. pylori	485	554	629

Table 1: Statistics of the eight datasets.

the remaining ones as training data. Due to the space limit, we only present the results of "Recovered number". The AUC results are in the supplemental material.

As shown in Figure 2, MATBoost generally achieves higher performance than other baselines. We observe that MATBoost recovers a significantly larger number of reactions than other baselines in datasets (a) – (c) and (f) – (h), and achieves competitive results with SHC in datasets (d) and (e). **This is because**: 1) MATBoost avoids directly performing inference in the incidence space that has size $O(2^M)$ by first inferring the adjacency matrix. This transforms an $O(2^M)$ problem into an $O(M^2)$ problem, which greatly reduces the problem size and eases the problem; 2) MATBoost leverages the powerful matrix factorization technique to perform inference in the adjacency space, which is able to well recover the unseen compound interactions given by ΔA^-; 3) MATBoost infers the incidence matrix by solving a constrained optimization problem, which recovers the reactions that best match the predicted compound interactions in adjacency space; 3) MATBoost takes advantage of a fixed-point iteration to boost performance and benefits from ensemble averaging. In all eight datasets, MATBoost is able to successfully recover over or nearly half of the missing reactions with the increasing of missing number. This high recall in predicting top reactions is very important because biologists usually can only verify top predictions due to resource and time limits.

In addition, we observe that SHC generally ranks second. This is an interesting result indicating that vertices and hyperlinks in a hypernetwork are sometimes interchangeable (by transposing the incidence matrix). We find that HKatz and HCN do not perform well. The phenomenon suggests that simply generalizing topological indices of links to hyperlinks by averaging may not be able to reflect the likelihood of a hyperlink. In terms of time complexity, MATBoost is also efficient. It takes only a few minutes to finish on the largest task.

Conclusions

In this paper, we consider the important task of predicting missing metabolic reactions for metabolic networks. By exploiting the problem structure, we model it as a hyperlink prediction problem. Hyperlink prediction is an interesting and challenging problem, however, relevant research are still limited. We proposed a novel algorithm for hyperlink prediction, called Matrix Boosting (MATBoost). It leverages the problem structure and implements an alternating C-M optimization framework in order to perform inference jointly in the incidence space and adjacency space. We have provided extensive evaluation results to demonstrate the unmatched performance of MATBoost. We compared our algorithm with six baselines. Experimental results on eight metabolic networks demonstrated that our MATBoost algorithm could recover over half of the missing reactions, which outperformed all existing methods by a large margin.

Hyperlink prediction may also be useful in many other application domains such as collaboration, authorship, and

social networks. We will study them in our future work.

References

[Bordbar et al. 2014] Bordbar, A.; Monk, J. M.; King, Z. A.; and Palsson, B. O. 2014. Constraint-based models predict metabolic and associated cellular functions. *Nature Reviews Genetics* 15(2):107–120.

[Bu et al. 2010] Bu, J.; Tan, S.; Chen, C.; Wang, C.; Wu, H.; Zhang, L.; and He, X. 2010. Music recommendation by unified hypergraph: combining social media information and music content. In *Proceedings of the 18th ACM international conference on Multimedia*, 391–400. ACM.

[Candès and Recht 2009] Candès, E. J., and Recht, B. 2009. Exact matrix completion via convex optimization. *Foundations of Computational mathematics* 9(6):717–772.

[Cao, Kong, and Yu 2014] Cao, B.; Kong, X.; and Yu, P. S. 2014. Collective prediction of multiple types of links in heterogeneous information networks. In *Data Mining (ICDM), 2014 IEEE International Conference on*, 50–59. IEEE.

[Chen et al. 2015] Chen, Z.; Chen, M.; Weinberger, K. Q.; and Zhang, W. 2015. Marginalized denoising for link prediction and multi-label learning. In *AAAI*, 1707–1713. Citeseer.

[Ermiş, Acar, and Cemgil 2015] Ermiş, B.; Acar, E.; and Cemgil, A. T. 2015. Link prediction in heterogeneous data via generalized coupled tensor factorization. *Data Mining and Knowledge Discovery* 29(1):203–236.

[Katz 1953] Katz, L. 1953. A new status index derived from sociometric analysis. *Psychometrika* 18(1):39–43.

[Khamsi and Kirk 2011] Khamsi, M. A., and Kirk, W. A. 2011. *An introduction to metric spaces and fixed point theory*, volume 53. John Wiley & Sons.

[Koren, Bell, and Volinsky 2009] Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *Computer* (8):30–37.

[Kumar, Dasika, and Maranas 2007] Kumar, V. S.; Dasika, M. S.; and Maranas, C. D. 2007. Optimization based automated curation of metabolic reconstructions. *BMC bioinformatics* 8(1):1.

[Liben-Nowell and Kleinberg 2007] Liben-Nowell, D., and Kleinberg, J. 2007. The link-prediction problem for social networks. *Journal of the American society for information science and technology* 58(7):1019–1031.

[Lü and Zhou 2010] Lü, L., and Zhou, T. 2010. Link prediction in weighted networks: The role of weak ties. *EPL (Europhysics Letters)* 89(1):18001.

[Lü and Zhou 2011] Lü, L., and Zhou, T. 2011. Link prediction in complex networks: A survey. *Physica A: Statistical Mechanics and its Applications* 390(6):1150–1170.

[Menon and Elkan 2011] Menon, A. K., and Elkan, C. 2011. Link prediction via matrix factorization. In *Machine Learning and Knowledge Discovery in Databases*. Springer. 437–452.

[Miller, Jordan, and Griffiths 2009] Miller, K.; Jordan, M. I.; and Griffiths, T. L. 2009. Nonparametric latent feature models for link prediction. In *Advances in neural information processing systems*, 1276–1284.

[Palais 2007] Palais, R. S. 2007. A simple proof of the banach contraction principle. *Journal of Fixed Point Theory and Applications* 2(2):221–223.

[Rendle et al. 2009] Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2009. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*, 452–461. AUAI Press.

[Rendle 2010] Rendle, S. 2010. Factorization machines. In *Data Mining (ICDM), 2010 IEEE 10th International Conference on*, 995–1000. IEEE.

[Rendle 2012] Rendle, S. 2012. Factorization machines with libfm. *ACM Transactions on Intelligent Systems and Technology (TIST)* 3(3):57.

[Shi and Malik 2000] Shi, J., and Malik, J. 2000. Normalized cuts and image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 22(8):888–905.

[Singh and Gordon 2008] Singh, A. P., and Gordon, G. J. 2008. Relational learning via collective matrix factorization. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, 650–658. ACM.

[Song, Meyer, and Tao 2015] Song, D.; Meyer, D. A.; and Tao, D. 2015. Top-k link recommendation in social networks. In *Data Mining (ICDM), 2015 IEEE International Conference on*, 389–398. IEEE.

[Sun, Ji, and Ye 2008] Sun, L.; Ji, S.; and Ye, J. 2008. Hypergraph spectral learning for multi-label classification. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, 668–676. ACM.

[Thiele and Palsson 2010] Thiele, I., and Palsson, B. Ø. 2010. A protocol for generating a high-quality genome-scale metabolic reconstruction. *Nature protocols* 5(1):93–121.

[Tian, Hwang, and Kuang 2009] Tian, Z.; Hwang, T.; and Kuang, R. 2009. A hypergraph-based learning algorithm for classifying gene expression and arraycgh data with prior knowledge. *Bioinformatics* 25(21):2831–2838.

[Wu et al. 2016] Wu, L.; Ge, Y.; Liu, Q.; Chen, E.; Long, B.; and Huang, Z. 2016. Modeling users preferences and social links in social networking services: A joint-evolving perspective. In *Thirtieth AAAI Conference on Artificial Intelligence*.

[Xu, Rockmore, and Kleinbaum 2013] Xu, Y.; Rockmore, D.; and Kleinbaum, A. M. 2013. Hyperlink prediction in hypernetworks using latent social features. In *Discovery Science*, 324–339. Springer.

[Yu, Tao, and Wang 2012] Yu, J.; Tao, D.; and Wang, M. 2012. Adaptive hypergraph learning and its application in image classification. *Image Processing, IEEE Transactions on* 21(7):3262–3272.

[Zhou, Huang, and Schölkopf 2005] Zhou, D.; Huang, J.; and Schölkopf, B. 2005. Learning from labeled and unlabeled data on a directed graph. In *Proceedings of the 22nd*

international conference on Machine learning, 1036–1043. ACM.

[Zhou, Huang, and Schölkopf 2006] Zhou, D.; Huang, J.; and Schölkopf, B. 2006. Learning with hypergraphs: Clustering, classification, and embedding. In *Advances in neural information processing systems*, 1601–1608.